

What is claimed is:

1        1. A method for use in a database system having plural storage modules,  
2 comprising:  
3                storing rows of a first table in a first storage module;  
4                storing rows of a second table in a second storage module;  
5                receiving a request to perform a join of the first and second tables;  
6                distributing, in response to the join request, rows and one or more columns  
7                of the rows of the first table from the first storage module to the second storage module;  
8                and  
9                sending row identifiers of the distributed rows with the distributed rows.

1        2. The method of claim 1, further comprising identifying the one or more  
2 columns of the first table that are part of one or more conditions of the join request.

1        3. The method of claim 2, wherein identifying the one or more columns that  
2 are part of the one or more conditions comprises identifying the one or more columns that  
3 are part of a join condition of the join request.

1        4. The method of claim 3, wherein identifying the one or more columns that  
2 are part of the one or more conditions further comprises identifying the one or more  
3 columns that are part of a residual condition of the join request.

1        5. The method of claim 3, wherein distributing the one or more columns does  
2 not comprise distributing columns that are part of a specified result list in the join request.

1        6. The method of claim 5, wherein distributing the one or more columns does  
2 not comprise distributing columns that are part of a subsequent operation specified in the  
3 join request.

1           7.       The method of claim 3, wherein distributing the one or more columns does  
2       not comprise distributing columns that are part of a subsequent operation specified in the  
3       join request.

1           8.       The method of claim 1, further comprising storing, in the second storage  
2       module, the distributed rows and one or more columns.

1           9.       The method of claim 8, wherein storing the distributed rows and one or  
2       more columns comprises storing in a spool table.

1           10.      The method of claim 9, wherein the second storage module is associated  
2       with an access module, the method further comprising performing a join, by the access  
3       module, of the spool table and the second table.

1           11.      The method of claim 10, wherein performing the join comprises  
2       identifying rows of the second table that satisfy one or more conditions of the join request  
3       and sending one or more columns of the identified rows from the second storage module  
4       to the first storage module.

1           12.      The method of claim 11, wherein sending the one or more columns of the  
2       identified rows comprises sending one or more columns of the second table that are part  
3       of the specified result list of the join request and that are part of a subsequent operation  
4       specified in the join request.

1           13.      The method of claim 12, wherein the first storage module is associated  
2       with one other access module, the method further comprising selecting, by the one other  
3       access module, rows of the first table corresponding to the identified rows of the second  
4       table and placing the selected rows of the first table and identified rows of the second  
5       table into a result table.

1           14. The method of claim 9, further comprising generating a temporary index  
2 based on the spool table, the temporary index to match a column of the second table to a  
3 row identifier in the spool table.

1           15. An article comprising at least one storage medium containing instructions  
2 executable in a database system having plural access modules to control access of plural  
3 storage modules, the instructions when executed causing the database system to:

4               store rows of a first table with a first access module;  
5               store rows of a second table with a second access module;  
6               receive a join request to join the first table and second table;  
7               identify one or more columns of the first table that are part of one or more  
8 conditions of the join request; and

9               distribute the identified one or more columns of the first table from the  
10 first access module to the second access module but not distributing columns of the first  
11 table that are part of a specified result list of the join request.

1           16. The article of claim 15, wherein the instructions when executed cause the  
2 database system to further:

3               distribute rows containing the one or more identified columns of the first  
4 table; and  
5               distribute row identifiers of the distributed rows with the distributed rows.

1           17. The article of claim 16, wherein the instructions when executed cause the  
2 database system to receive, by the first access module, rows of the second table that  
3 satisfy the one or more join conditions of the join request.

1           18. The article of claim 17, wherein the instructions when executed cause the  
2 database system to receive the rows of the second table by receiving rows containing one  
3 or more columns of the second table that are part of the specified result list in the join  
4 request.

1           19. The article of claim 18, wherein the instructions when executed cause the  
2 database system to receive the rows of the second table containing one or more further  
3 columns that are part of a subsequent operation specified in the join request.

1           20. The article of claim 18, wherein the instructions when executed cause the  
2 database system to further not distribute columns of the first table that are part of a  
3 subsequent operation specified in the join request.

1           21. The article of claim 17, wherein the instructions when executed cause the  
2 database system to further:

3                   receive the row identifiers of the first table along with the rows of the  
4 second table;  
5                   retrieve rows from the first table using the row identifiers; and  
6                   storing rows of first and second tables in a result table.

1           22. A database system comprising:  
2                   a plurality of storage modules, with a first storage module storing rows of  
3 a first table and a second storage module storing rows of a second table;  
4                   a plurality of access modules adapted to manage access of respective  
5 storage modules, a first access module corresponding to the first storage module, and a  
6 second access module corresponding to the second storage module; and  
7                   the first access module adapted to distribute rows of the first table to the  
8 second access module in response to a join request, the first access module adapted to  
9 further distribute row identifiers of the distributed rows with the distributed rows.

1           23. The database system of claim 22, wherein the distributed rows contain one  
2 or more columns that are part of one or more join conditions of the join request but do not  
3 contain one or more columns that are in a specified result list of the join request.

1        24. The database system of claim 23, wherein the distributed rows do not  
2 contain one or more columns that are part of a subsequent operation specified in the join  
3 request.

1        25. The database system of claim 22, wherein each row of the first table is  
2 associated with a primary index, and wherein each row identifier comprises a hash code  
3 of the primary index and a uniqueness value.

1        26. The database system of claim 25, wherein the first table is hash partitioned  
2 and hash ordered

1        27. The database system of claim 22, wherein the first table is hash partitioned  
2 and value ordered, and wherein each row of the first table is associated with a primary  
3 index, and wherein each row identifier comprises a hash code of the primary index and a  
4 field used for value ordering.

1        28. The database system of claim 22, wherein the first table is value  
2 partitioned and hash ordered, and wherein each row of the first table is associated with a  
3 primary index, and wherein each row identifier comprises a hash code of the primary  
4 index, a uniqueness field, and a field used for value partitioning.

1        29. The database system of claim 22, wherein the first table is value  
2 partitioned and value ordered, and wherein each row identifier comprises at least one  
3 field used for one of value partitioning and for value ordering.

30. A database system comprising:

- a plurality of storage modules, with a first storage module storing rows of a first table and a second storage module storing rows of a second table;
- a plurality of access modules adapted to manage access of respective storage modules, a first access module corresponding to the first storage module, and a second access module corresponding to the second storage module; and
- the first access module adapted to distribute rows of the first table of the second access module in response to a join request, the first access module adapted to further distribute row identifiers with the distributed rows,
- the first access module adapted to further distribute columns of the first table that are part of one or more join conditions of the join request but to not distribute columns of the first table that are part of a specified result list in the join request and subsequent operation of the join request.

Digitized by srujanika@gmail.com